arXiv:1608.05617v1 [cs.PL] 19 Aug 2016

CurryCheck: Checking Properties of Curry Programs

Michael Hanus

Institut für Informatik, CAU Kiel, D-24098 Kiel, Germany
mh@informatik.uni-kiel.de

Abstract. We present CurryCheck, a tool to automate the testing of programs written in the functional logic programming language Curry. CurryCheck executes unit tests as well as property tests which are parameterized over one or more arguments. In the latter case, CurryCheck tests these properties by systematically enumerating test cases so that, for smaller finite domains, CurryCheck can actually prove properties. Unit tests and properties can be defined in a Curry module without being exported. Thus, they are also useful to document the intended semantics of the source code. Furthermore, CurryCheck also supports the automated checking of specifications and contracts occurring in source programs. Hence, CurryCheck is a useful tool that contributes to the property- and specification-based development of reliable and well tested declarative programs.

1 Motivation

Testing is an important step to get confidence in the functionality of a program. The advantage of testing compared to program verification is its potential for automation. If we do not execute test cases only manually for some inputs but encode them as input to test frameworks, we can automatically run and repeat them when the software is further developed, which is also known as regression testing.

A difficulty in testing is to find appropriate inputs for the individual tests. For this purpose, property-based testing has been proposed, well known in the functional language Haskell with the QuickCheck tool [16]. Basically, properties are predicates parameterized over one or more arguments. QuickCheck automates the test execution by applying properties to randomly generated test inputs. Since this idea is particularly reasonable for declarative languages, it is been adapted in different forms to functional and logic programming languages. For instance, SmallCheck [33] and GAST [28] focus on a systematic enumeration of test inputs for functional programs, PropEr [30] adapts ideas of QuickCheck to the concurrent functional language Erlang, PrologCheck [1] transfers and extends ideas of QuickCheck to Prolog, and EasyCheck [15] exploits functional logic programming features to property-based testing of Curry programs.

CurryCheck follows the same ideas. Actually, it is based on EasyCheck to define properties. However, CurryCheck is intended as a comprehensive tool to simplify the automation of test execution. To use CurryCheck, properties are interspersed into the program as top-level definitions. Thus, properties are used to document the intended semantics of the source code, which also supports test-driven program development known as "extreme programming." When CurryCheck is applied to a (set of) Curry modules, it extracts all properties, generates a program to test these properties, executes this generated program, and reports any errors. Furthermore, CurryCheck also analyzes possible contracts [8] provided in source programs and generates properties to test these contracts. Thanks to this automation, CurryCheck is a useful tool for continuous integration and deployment processes. Actually, it is used for this purpose in the Curry implementations PAKCS [24] and KiCS2 [14].

In this paper we present the ideas and usage of CurryCheck. After a review of the main features of Curry in the next section, we introduce properties in Sect. 3 and explain how they

are tested in Sect. 4. The support of CurryCheck to define test inputs is presented in Sect. 5. CurryCheck's support for contract checking is described in Sect. 6. Some initial features of CurryCheck to combine testing and verification are sketched in Sect. 7. We report about our practical experience with CurryCheck in Sect. 8 before we compare CurryCheck to some related tools and conclude.

2 Functional Logic Programming and Curry

Functional logic languages [6,23] integrate the most important features of functional and logic languages in order to provide a variety of programming concepts. They support functional concepts like higher-order functions and lazy evaluation as well as logic programming concepts like non-deterministic search and computing with partial information. This combination led to new design patterns [7] as well as better abstractions for application programming. The declarative multi-paradigm language Curry [20] is a modern functional logic language with advanced programming concepts. In the following, we briefly review some features of Curry relevant for this paper. More details can be found in recent surveys on functional logic programming [6,23] and in the language report [25].

The syntax of Curry is close to Haskell [31]. In addition to Haskell, Curry allows *free* (*logic*) *variables* in rules and initial expressions. Function calls with free variables are evaluated by a possibly non-deterministic instantiation of demanded arguments.

Example 1. The following simple program shows the functional and logic features of Curry. It defines the well-known list concatenation and an operation that returns some element of a list having at least two occurrences:

```
(++) :: [a] → [a] → [a]
[]     ++ ys = ys
(x:xs) ++ ys = x : (xs ++ ys)
```

```
someDup :: [a] → a
someDup xs | xs == _++[x]++_++[x]++_
           = x    where x free
```

The (optional) type declaration ("`::`") of the operation "`++`" specifies that "`++`" takes two lists as input and produces an output list, where all list elements are of the same (unspecified) type. Since "`++`" can be called with free variables in arguments, the condition in the rule of `someDup` is solved by instantiating `x` and the anonymous free variables "`_`" to appropriate values before reducing the function calls. This corresponds to narrowing [34,32], but Curry narrows with possibly non-most-general unifiers to ensure the optimality of computations [2].

Note that `someDup` is a *non-deterministic operation* since it might deliver more than one result for a given argument, e.g., the evaluation of `someDup [1,2,2,1]` yields the values `1` and `2`. Non-deterministic operations, which can formally be interpreted as mappings from values into sets of values [19], are an important feature of contemporary functional logic languages. Hence, Curry has also a predefined *choice* operation:

```
x ? _  =  x
_ ? y  =  y
```

Thus, the expression "`0 ? 1`" evaluates to `0` and `1` with the value non-deterministically chosen.

Functional patterns [3] are useful to define some operations more easily. A functional pattern is a pattern occurring in an argument of the left-hand side of a rule containing defined operations (and not only data constructors and variables). Such a pattern abbreviates the set of all standard patterns to which the functional pattern can be evaluated (by narrowing). For instance, we can rewrite the definition of `someDup` as

```
someDup (_++[x]++_++[x]++_) = x
```

Functional patterns are a powerful feature to express arbitrary selections in tree structures, e.g., in XML documents [22]. Details about their semantics and a constructive implementation of functional patterns by a demand-driven unification procedure can be found in [3].

Curry has also features which are useful for application programming, like *set functions* [5] to encapsulate non-deterministic computations, *default rules* [9] to deal with partially specified operations and negation, and standard features from functional programming, like modules or monadic I/O [36]. Other features are explained when they are used in the following.

3 Properties

In this section we briefly discuss which kind of program properties to be tested are supported by CurryCheck. Since CurryCheck extends the functionality of EasyCheck [15], it supports all kinds of EasyCheck's properties which we review first.

Properties are defined top-level entities with a distinct type (see below). Thus, their syntax and type-correctness can be checked by the standard front end of any Curry system. Properties do not require a specific naming convention but CurryCheck recognizes them by their type. Moreover, the name and position of the property in the source file are used by CurryCheck to identify properties when errors are reported.

For instance, consider the list concatenation operation "`++`" defined in Example 1. Before discussing general properties, we define some unit tests for fixed arguments, like

```
concNull12  =  []   ++ [1,2] -=- [1,2]
concCurry   =  "Cu" ++ "rry" -=- "Curry"
```

The infix operator "`-=-`" specifies a test which is successful if both sides have single values which are identical (we will later see tests for non-deterministic operations). Since the expressions can be of any type (of course, the two arguments must be of the same type), the operator is polymorphic and has the type

```
(-=-) :: a → a → Prop
```

Hence, all entities defined above have type `Prop`.

The power of CurryCheck and similar property-based test frameworks comes from the fact that we can also test properties which are parameterized over some input data. For instance, we can check whether the concatenation operation is associative by:

```
concIsAssociative xs ys zs = (xs++ys)++zs -=- xs++(ys++zs)
```

This property is parameterized over three input values `xs`, `ys`, and `zs`. To test this property, CurryCheck guesses values for these parameters (see below for more details) and tests the property for these values:

```
concIsAssociative_ON_BASETYPE (module ConcDup, line 18):
 OK, passed 100 tests.
```

Indicated by the suffix `_ON_BASETYPE`, we see another feature of CurryCheck. If properties are polymorphic in their input values (the above property has type `[a]` → `[a]` → `[a]` → `Prop`), CurryCheck specializes the type to some base type, since there is no concrete value of a polymorphic type (and EasyCheck would fail on such properties). As a default, CurryCheck uses the predefined type `Ordering` having the three values `LT`, `EQ`, `GT` (another more involved method to instantiate polymorphic types in purely functional programs can be found in [12]). This default type can be changed to other base types, like `Bool`, `Int`, or `Char`, with a command-line option. One could also provide an explicit type declaration for the property. For instance, we can test the commutativity of the list concatenation on lists of integers by the property

```
concIsCommutative :: [Int] → [Int] → Prop
concIsCommutative xs ys = (xs ++ ys) -=- (ys ++ xs)
```

Of course, this property does not hold so that CurryCheck reports an error together with a counter-example:

```
...
concIsCommutative (module ConcDup, line 20) failed
Falsified by 8th test.
Arguments: [-1] [-3]
Results:   ([-1,-3],[-3,-1])
```

Note that the arguments of a test are ordinary expressions so that one can use any defined operation in the tests. For instance, we can (sucessfully) check whether the list concatenation is the addition on their lengths:

```
concAddLengths xs ys = length xs + length ys -=- length (xs++ys)
```

Since Curry covers also logic programming features, CurryCheck supports the testing of non-deterministic properties. For instance, one can check whether an expression reduces to some given value with the operator is "`~>`":

```
someDup1 = someDup [1,2,1,2] ~> 1
```

Another important operator is "`<~>`" which denotes a test which succeeds if both arguments have the same set of values. We can write unit tests by enumerating all expected values with the choice operator "`?`":

```
someDup12 = someDup [1,2,1,2,1] <~> (1?2)
```

It should be noted that the operator "`<~>`" really compares sets and not multi-sets: Although the evaluation of `someDup [1,2,1,2,1]` returns the value `1` three times in a typical Curry system, the property `someDup12` holds. This is intended since CurryCheck tests declarative properties which are independent of specific compiler optimizations (this is in contrast to PrologCheck which tests operational properties like multiplicity of answers and modes [1]).

As another example, consider the following definition of a permutation of a list by exploiting a functional pattern to select some element in the argument list:

```
perm (xs++[x]++ys) = x : perm (xs++ys)
perm []            = []
```

An important property of a permutation is that the length of the list is not changed. Hence, we check it by the property

```
permLength xs = length (perm xs) <~> length xs
```

Note that the use of "`<~>`" (instead of "`-=-`") is relevant since non-deterministic values are compared. Actually, the left argument evaluates to many (identical) values.

We might also want to check whether our definition of `perm` computes the correct number of solutions. Since we know that a list of length n has $n!$ permutations, we write the following property, where `fac` is the factorial function and the property x `#` n is satisfied if x has n different values:

```
permCount :: [Int] → Prop
permCount xs = perm xs # fac (length xs)
```

However, this test will be falsified with the test input `[1,1]`, since this list has only one permuted value (actually, both computed values are identical). We can obtain a correct property if we add the condition that all elements in the input list `xs` are pairwise different. For this purpose, we use a *conditional property*: the property b `==>` p is satisfied if p is satisfied for all

values where b evaluates to `True`. If the predicate `allDifferent` is satisfied iff its argument list does not contain duplicated elements, then we can reformulate our property as follows:

```
permCount xs = allDifferent xs ==> perm xs # fac (length xs)
```

Furthermore, we want to check the existence of distinguished permutations. For this purpose, consider a predicate to check whether a list is sorted:

```
sorted :: [Int] → Bool
sorted []       = True
sorted [_]      = True
sorted (x:y:zs) = x<=y && sorted (y:zs)
```

Then we can check whether there are sorted permutations (the property "`eventually` x" is satisfied if some value of x is `True`):

```
permIsEventuallySorted :: [Int] → Prop
permIsEventuallySorted xs = eventually (sorted (perm xs))
```

Property-based testing is appropriate for declarative languages since the absence of side effects allows the execution of tests on any number of test data without influencing the individual tests. Nevertheless, real programming languages have to deal with the real world so that they support also I/O operations. Clearly, such operations should also be tested. Although there are methods to test monadic code [17], Curry supports only I/O monadic operations where testing with arbitrary data seems not reasonable. Therefore, CurryCheck supports only non-parameterized unit tests for I/O operations. For instance, the test (a `` `returns` `` x) is satisfied if the I/O action a returns the value x. For instance, we can test whether writing a file and reading it yields the same contents:

```
writeReadFile = (writeFile "TEST" "Hello" >> readFile "TEST")
                `returns` "Hello"
```

Since CurryCheck executes the tests written in a source program in their textual order, one can write also several I/O tests whose side effects depend on each other. For instance, we can split the previous I/O test into two consecutive tests:

```
writeTestFile = (writeFile "TEST" "Hello") `returns` ()
readTestFile  = (readFile "TEST")          `returns` "Hello"
```

4 Testing Properties

After having seen several methods to define properties, we sketch in this section how they are actually tested. Our motivation for the development of CurryCheck is manifold:

1. Properties are an executable documentation for the intended semantics of operations.
2. Properties increase the confidence in the quality of the developed software.
3. Properties can be used for software verification by proving their validity for all possible input data.

The first point is supported by interspersing properties into the source code of the program instead of putting them into separate files. Thus, properties play the same role as comments or type annotations: they document the intended semantics. Hence, they can be extracted and put into the program documentation by automatic documentation tools [21,26]. In order to avoid that properties influence the interface of a module, they do not need to be exported. As an example, consider the following simple module defining the classical list reverse operation (the imported module `Test.EasyCheck` contains the definitions of the property combinators introduced in Sect. 3):

```
module Rev(rev) where
import Test.EasyCheck
rev :: [a] → [a]
rev []     = []
rev (x:xs) = rev xs ++ [x]
revLength  xs = length (rev xs) -=- length xs
revRevIsId xs = rev (rev xs) -=- xs
```

We can run all tests of this module by invoking the CurryCheck executable with the name of the module:[1]

```
> currycheck Rev
Analyzing module 'Rev'...
...
Executing all tests...
revLength_ON_BASETYPE (module Rev, line 9):
 OK, passed 100 tests.
revRevIsId_ON_BASETYPE (module Rev, line 10):
 OK, passed 100 tests.
```

Although module `Rev` only exports the operation `rev`, all properties defined in the top-level of `Rev` are passed to the underlying EasyCheck library for testing. For this purpose, CurryCheck creates a copy of this module where all entities are exported (note that the original module cannot be modified since it might be imported to other modules to be tested). For each property a corresponding call to an operation of EasyCheck is generated which actually performs the generation of test data, runs the test, and collects all results which are passed back to CurryCheck. Furthermore, polymorphic properties are not checked but a corresponding new property on the default base type is generated which calls the polymorphic property. For instance, if the default base type is `Int`, CurryCheck generates the new property

```
revRevIsId_ON_BASETYPE :: [Int] → Prop
revRevIsId_ON_BASETYPE = revRevIsId
```

which is actually checked instead of `revRevIsId`. Note that it might lead to a failure if the type of `revRevIsId` is directly specialized, since the polymorphic property `revRevIsId` might be used in other property definitions with a different specialized type.

After these preparations, EasyCheck tests the properties by generating test data as described in [15]. EasyCheck does not use random generators like QuickCheck or PrologCheck, but it exploits functional logic programming features to enumerate possible input values. Since logic variables are equivalent to non-deterministic generators [4], one can evaluate a logic variable of a particular type in order to get all values of this type in a non-deterministic manner. For instance, if we evaluate the Boolean variable `b::Bool`, we obtain the values `False` and `True` as results. Similarly, for `bs::[Bool]` we obtain values like `[]`, `[False]`, `[True]`, `[False,False]`, etc. In order to select a finite amount of these infinite values, one can use Curry's feature for encapsulated search to collect all non-deterministic results in a tree structure, traverse the tree with some strategy and return the result of the traversal into a list. If one selects only a finite amount of this list, the lazy evaluation strategy of Curry ensures a finite computation even if the tree is infinite. Based on these features, the EasyCheck library contains an operation

```
valuesOf :: a → [a]
```

[1] One can also provide several module names so that they are tested at once. Furthermore, CurryCheck has various options to influence the number of test cases, default types for polymorphic properties, etc.

which computes the list of all values of the given argument according to a fixed strategy (in the current implementation by randomized level diagonalization [15]). Hence, we can get 20 values for a list of integers by

```
...> take 20 (valuesOf (_::[Int]))
[[],[-1],[-3],[0],[1],[-1,0],[-2],[0,0],[3],[-1,1],[-3,0],[0,1],
[2],[-1,-1],[-5],[0,-1],[5],[-1,2],[-9],[0,2]]
```

It should be noted that `valuesOf` enumerates all values of the given type completely and without duplicates.[2] Hence, if the set of possible input values is finite, it is ensured that all of them are tested if sufficiently many tests are performed. In this case, the property is also verified (where QuickCheck or PrologCheck does not give such guarantees). For instance, consider the De Morgan law from Boolean algebra:

```
negOr b1 b2 = not (b1 || b2) -=- not b1 && not b2
```

This property is proved by CurryCheck after four tests with all possible input values, and the output of CurryCheck indicates that the testing was exhaustive:

```
negOr (module BoolTest, line 4):
 Passed all available tests: 4 tests.
```

5 User-Defined Test Data

Due to the use of functional logic features to generate test data, one can write properties not only on predefined data types but also on user-defined data types. For instance, consider the following definition of general polymorphic trees:

```
data Tree a = Leaf a | Node [Tree a]
```

We define operations to compute the leaves of a tree and mirroring a tree:

```
leaves (Leaf x)  = [x]
leaves (Node ts) = concatMap leaves ts

mirror (Leaf x)  = Leaf x
mirror (Node ts) = Node (reverse (map mirror ts))
```

The following properties should increase our confidence in the correctness of the implementation:

```
doubleMirrorIsId t = mirror (mirror t) -=- t

leavesOfMirrorAreReversed t = leaves t -=- reverse (leaves (mirror t))
```

CurryCheck successfully tests these properties without providing any further information about how to generate test data. However, in some cases it might be desirable to define our own test data since the generated structures are not appropriate for testing. For instance, if we test algorithms working on balanced search trees, we need correctly balanced search trees as test data. As a naive approach, we can limit the tests to correct test inputs by using conditional properties. As a simple example, consider the following operation that adds all numbers from 1 to a given limit:

```
sumUp n = if n==1 then 1 else n + sumUp (n-1)
```

Since there is also a simple formula to compute this sum, we can check it:

```
sumUpIsCorrect n = n>0 ==> sumUp n -=- n * (n+1) `div` 2
```

[2] In order to get an idea of the distribution of the generated test data, CurryCheck also provides property combinators `collect` and `classify` known from QuickCheck.

Note that the condition is important since `sumUp` diverges on non-positive numbers. As a result, CurryCheck tests this property by enumerating integers and dropping tests with non-positive numbers. While this works well, since CurryCheck performs a fairly good distribution between positive and negative numbers, this approach might have a serious drawback if the proportion of correct test data is small. In the case of balanced search trees, there are many more unbalanced trees than balanced search trees. This has the effect that CurryCheck generates many test data and drops it so that it does not make much progress. Actually, CurryCheck has an upper limit for dropping test data in the conditional operator in order to avoid spending too much work on generating unusable test data. For instance, if we want to test the above property `revRevIsId` on long input lists, we could define it as follows:

```
revRevIsIdLong :: [Int] → Prop
revRevIsIdLong xs = length xs > 100 ==> rev (rev xs) -=- xs
```

Since there are a huge number of integer lists with a length smaller than 100, CurryCheck does not find any test case (with a default limit of dropping at most 10,000 incorrect test data values):

```
revRevIsIdLong (module Rev, line 13):
 Arguments exhausted after 0 test.
```

This shows that the fully automatic generation of test data is not always appropriate. Therefore, CurryCheck provides some combinators to explicitly define test data (more advanced enumeration combinators in the context of Scala are discussed in [29]).

To show the method for test data generation in more detail, we have to review Curry's methods to encapsulate non-deterministic computations. For this purpose, Curry defines the following structure to represent the results of a non-deterministic computation [13]:

```
data SearchTree a = Value a | Fail | Or (SearchTree a) (SearchTree a)
```

`(Value v)` and `Fail` represent a single value or a failure (i.e., no value), respectively, and `(Or t1 t2)` represents a non-deterministic choice between two search trees `t1` and `t2`. Furthermore, there is a primitive search operator

```
someSearchTree :: a → SearchTree a
```

which yields a search tree for an expression. For instance, `someSearchTree (0?1)` evaluates to the search tree

```
Or (Value 0) (Value 1)
```

The expression

```
someSearchTree (id $## (_::[Bool]))
```

(where "`$##`" is an infix application operator which evaluates the right argument to ground normal form before applying the left argument to it) yields an (infinite) search tree of all Boolean lists:

```
(Or (Value []) (Or (Or (Or (Value [False]) ...) (Or ...)) ...))
```

Basically, EasyCheck defines various strategies to traverse such search trees (see [15] for details) in order to enumerate test data. Hence, if we want to define our own test data, we have to define an operation that generates a search tree containing the test data in `Value` leaves. Although this is not difficult for simple data types, it could be demanding for polymorphic types where generators for the polymorphic arguments must be weaved with the generators for the main data structure. To simplify this task, CurryCheck offers a family of combinators `genCons`n where each combinator takes an n-ary constructor function and n generators as arguments and produces a search tree for this constructor and all combinations of generated arguments. Hence, these combinators have the type

`genCons`n `::` $(a_1 \to \cdots \to a_n \to a)$ $\to$ `SearchTree` a_1 $\to \cdots \to$ `SearchTree` a_n $\to$ `SearchTree` a

Furthermore, there is an infix combinator "`|||`" to combine two search trees. For instance, consider the straightforward definition of Peano numbers:

```
data Nat = Z | S Nat
```

Then we can define a search tree generator for this type as follows:

```
genNat :: SearchTree Nat
genNat = genCons0 Z ||| genCons1 S genNat
```

Similarly, we can define a search tree generator for polymorphic trees which takes a search tree for the argument type as a parameter (where `genList` denotes the corresponding generator for list values):

```
genTree :: SearchTree a → SearchTree (Tree a)
genTree ta = genCons1 Leaf ta ||| genCons1 Node (genList (genTree ta))
```

The explicit definition of value generators is reasonable when only a subset of all values should be used for testing. For instance, `sumUpIsCorrect` should be testest with positive numbers only. Hence, we define a generator for positive numbers:

```
genPos = genCons0 1 ||| genCons1 (+1) genPos
```

Since these numbers are slowly increasing, i.e., the search tree is actually degenerated to a list, we can also use the following definition to obtain a more balanced search tree:

```
genPos = genCons0 1 ||| genCons1 (\n → 2*(n+1)) genPos
                    ||| genCons1 (\n → 2*n+1)   genPos
```

In order to test properties with user-defined data, CurryCheck provides the property combinator

```
forAll :: [a] → (a → Prop) → Prop
```

which is satisfied if the parameterized property given as the second argument is satisfied for all values of the first argument list. Since there is also a library operation

```
valuesOfSearchTree :: SearchTree a → [a]
```

(actually, the operation `valuesOf` introduced in Sect. 4 is defined via this operation) to enumerate all values of a search tree, we can redefine the property `sumUpIsCorrect` as follows:

```
sumUpIsCorrect = forAll (valuesOfSearchTree genPos)
                        (\n → sumUp n -=- n*(n+1) `div` 2)
```

Using this technique, we could also define finite tests for potentially infinite structures, e.g., one can easily define tree generators that generate all trees up to a particular depth.

Finally, we show the implementation of the combinators to generate search trees. The definition of "`|||`" and `genCons0` is straightforward:

```
x ||| y = Or x y
genCons0 v = Value v
```

To define the further combinators like `genCons1`, we have to replace in a given search tree (for the argument) the `Value` nodes by new nodes where the constructor operation is applied to the given value. This task is done by the following auxiliary operation:[3]

```
updateValues :: SearchTree a → (a → SearchTree b) → SearchTree b
updateValues (Value a)  f = f a
updateValues Fail       f = Fail
updateValues (Or t1 t2) f = Or (updateValues t1 f) (updateValues t2 f)
```

[3] This operation is similar to the monadic bind operation in Haskell's `MonadPlus`, but we use this definition due to the lack of type classes in the current language definiton of Curry.

The definition of the remaining combinators is now easy (we only show the first two ones):

```
genCons1 c gena = updateValues gena (\a → Value (c a))
genCons2 c gena1 gena2 =
  updateValues gena1 (\a1 → updateValues gena2 (\a2 → Value (c a1 a2)))
```

6 Contract and Specification Testing

As discussed in detail in [8], the distinctive features of Curry (e.g., non-deterministic operations, demand-driven evaluation, functional patterns, set functions) support writing high-level specifications as well as efficient implementations for a given problem in the same programming language. When applying this idea, Curry can be used as a wide-spectrum language [11] for software development. If a specification or contract is provided for some function, one can exploit this information to support run-time assertion checking with these specifications and contracts. As an additional use of this information, CurryCheck automatically generates properties to test the given specifications and contracts, which is described in the following.

According to the notation proposed in [8], a *specification* for an operation f is an operation `f'spec` of the same type as f. A *contract* consists of a pre- and a postcondition, where the precondition could be omitted. When provided, a *precondition* for an operation f of type $\tau \to \tau'$ is an operation

```
f'pre :: τ → Bool
```

putting demands on allowed arguments, whereas a *postcondition* for f is an operation

```
f'post :: τ → τ' → Bool
```

which relates input and output values (the generalization to operations with more than one argument is straightforward). A specification should precisely describe the meaning of an operation, i.e., the declarative meaning of the specification and the implementation of an operation should be equivalent. In contrast, a contract is a partial specification, e.g., all results computed by the implementation should satisfy the postcondition.

As a concrete example, consider the problem of sorting a list. The specification defines a sorted version of a given list as a permutation of the input which is sorted. Exploiting the operations introduced in Sect.3, we define the following specification for the operation `sort`:

```
sort'spec :: [Int] → [Int]
sort'spec xs | ys == perm xs && sorted ys = ys  where ys free
```

A postcondition, which is easier to check, states that the input and output lists should have the same length:

```
sort'post :: [Int] → [Int] → Bool
sort'post xs ys = length xs == length ys
```

To provide a concrete implementation, we implement the quicksort algorithm as follows:

```
sort :: [Int] → [Int]
sort []     = []
sort (x:xs) = sort (filter (<x) xs) ++ [x] ++ sort (filter (>x) xs)
```

Note that specifications and contracts are optional. However, if they are included in a module processed with CurryCheck, CurryCheck automatically generates and checks properties that relate the specification and contract to the implementation. For instance, an implementation satisfies a specification if both yield the same values, and a postcondition is satisfied if each value computed for some input satisfies the postcondition relation between input and output.

For our example, CurryCheck generates the following properties (if there are also preconditions for some operation, these preconditions are used to restrict the test cases via the condition operater "==>"):[4]

```
sortSatisfiesSpecification :: [Int] → Prop
sortSatisfiesSpecification x = sort x <~> sort'spec x

sortSatisfiesPostCondition :: [Int] → Prop
sortSatisfiesPostCondition x = always (sort'post x (sort x))
```

With CurryCheck, the framework of [8] becomes more useful since contracts are not only used as run-time assertions in concrete computations, but many possible computations are checked with various test data. For instance, CurryCheck reports that the above implementation of `sort` is incorrect for the example input `[1,1]` (as the careful reader might have already noticed). When reporting the error, the module and source code line number of the erroneous operation is shown so that the programmer can easily spot the problem.

Another kind of contracts taken into account by CurryCheck are determinism annotations. An operation that yields always identical results (maybe multiple times) on identical argument values can be annotated as "deterministic" by adding `DET` to the result type of its type signature. For instance, the following operation tests whether a list represents a set, i.e., has no duplicate elements (the definition exploits functional patterns [3] as well as default rules [9]):

```
isSet :: [a] →DET Bool
isSet (_++[x]++_++[x]++_) = False
isSet'default _           = True
```

The determinism annotation "→`DET`" has the effect that at most one result is computed for a given input, e.g., a single value `False` is returned for the call `isSet [1,3,1,3,1]`, although the first rule can be applied multiple times to this call. Thus, after computing a first value, all attempts to compute further values are ignored. In order to ensure that this does not destroy completeness, i.e., it behaves like "green cuts" in Prolog, such operations must be deterministic from a semantical point of view. CurryCheck tests this property by generating a property for each `DET`-annotated operation that expresses that there is at most one value for each input. For instance, for `isSet`, the `DET` annotation is removed and the property

```
isSetIsDeterministic x1 = isSet x1 #< 2
```

is added by CurryCheck, where "e `#<` n" is satisfied if the *set* of all values of e contains less than n elements.

7 Combining Testing and Verification

The objective of CurryCheck is to increase the confidence in the reliability of Curry programs. Testing with a lot of input data is one important step but, in case of input data types with infinite values, it can only show possible errors but not the absence of them. In order to support the latter, CurryCheck has also some (preliminary) support to include the verification of program properties. For this purpose, a programmer might prove properties stated in a source program. Since there are many different possibilities to prove such properties, ranging from manual proofs to interactive proof assistants and fully automatic provers, CurryCheck does not enforce a particular proof technique. Instead, CurryCheck trusts the programmer and uses a naming convention for files containing proofs: if there is a file with name `proof-`t`.*`, CurryCheck assumes that this file contains a valid proof for property t. For instance, the following property states that sorting a list does not change its length:

[4] The property "`always` x" is satisfied if all values of x are `True`.

```
sortlength xs = length (sort xs) <~> length xs
```

If there is a file `proof-sortlength.agda`, containing a proof for the above property ([10] addresses techniques how to prove such properties in the dependently typed language Agda), CurryCheck considers this property as valid and does not check it. Moreover, it uses it to simplify other properties to be tested. In our case, the property `sortSatisfiesPostCondition` of the previous section can be simplified to `always True` so that it does not need to be tested. Similarly, a determinism annotation for operation f is not tested if there is a proof file f`IsDeterministic.*`.

Since program verification is a notoriously difficult task, a mixture of different techniques is required. For instance, [27] discusses techniques to use the Isabelle/HOL proof assistant to verify purely functional properties inspired by QuickCheck. [10] describes a method to prove non-deterministic computations by translating Curry programs into Agda programs. Since these proofs can be verified by the Agda compiler, CurryCheck can test the validity of a given proof file by simply invoking the Agda compiler. Some purely functional properties can be proved in a fully automatic way. For instance, the properties

```
concLength xs = length (xs ++ ys) -=- length xs + length ys
revLength  xs = length (rev xs) -=- length xs
```

can be proved by the SMT solver Alt-Ergo. To support the use of such solvers, we have started the development of tools to automatically translate Curry programs into the syntax of Agda and other proof systems. We omit more details since this is outside the scope of this paper.

8 Practical Experience

The implementation of CurryCheck is available with the (Prolog-based) Curry implementation PAKCS [24] (since version 1.14.0) and the (Haskell-based) Curry implementation KiCS2 [14] (since version 0.5.0). The implementation exploits meta-programming features available in these systems to parse programs and transform them into new programs as described in the previous sections.

Although we could show in this paper only simple examples, we would like to remark that CurryCheck is successfully applied in a larger context. CurryCheck is regularly used for automatic regression testing during continuous integration and nightly builds of PAKCS and KiCS2. Currently, approximately 500 properties (the number is continuously growing) are regularly used to test the libraries and tools of these systems. Our practical experience is quite promising. After the development and use of CurryCheck, we found a bug in the implementation of the prelude operations `quot` and `rem` w.r.t. negative numbers and free variables which was undetected for a couple of years. Although the bug was easy to fix, the definition of a general property relating both operations and testing it with all smaller values was essential for its detection.

The run time of CurryCheck mainly depends on the specific properties to be tested. The initial translation phase, which extracts properties, contracts, and specifications from a given module and transforms them into executable tests, is a straightforward compilation process. The run time of the subsequent test execution phase depends on the number of test cases and the time needed to evaluate each property. The functional logic programming technique to generate test data described in Sect. 4 (i.e,. collecting all non-deterministic results of evaluating a logic variable) is reasonable in practice. For instance, KiCS2 needs 0.6 seconds to test a trivial property on a list of integers with 10,000 test cases computed by the randomized level diag-

onalization strategy described in [15] (on a Linux machine with Intel Core i7-4790/3.60Ghz and 8GiB of memory).

9 Related Work

Since testing is an important part of the software development process, there is a vast literature on this topic. In the following, we compare our approach to testing, in particular, property-based testing, in declarative languages. We already mentioned QuickCheck [16] which was influential in this area and followed by other property-testing systems for functional languages, like GAST [28] or SmallCheck [33]. The same idea has also been transferred to other languages like PropEr [30] for Erlang and PrologCheck [1] for Prolog. In contrast to CurryCheck, most of these systems (except for SmallCheck) are based on randomly generating test data so that they do not provide guarantees for a complete enumeration if the sets of input values are finite, i.e., they cannot verify properties. PropEr also supports contract checking but these function contracts are limited to type specifications. PrologCheck could also check operational aspects likes modes or multiplicity of answers, whereas our properties are oriented towards declarative aspects, i.e., the input/output relation of values.

Closely related to CurryCheck is EasyCheck [15] since it can be seen as our back end. EasyCheck is the only property-based test tool covering functional and logic aspects but it is more limited than CurryCheck. EasyCheck does not support polymorphic properties, I/O properties, or combinators for user-defined generation of test data. This has been added in CurryCheck together with a full automation of the test process in order to obtain an easily usable tool. Moreover, CurryCheck expands the use of automatic testing by using it for contract and specification checking, where functional logic programming has been shown to be an appropriate framework [8], and combines it with static verification techniques.

10 Conclusion

We have presented CurryCheck, the first fully automatic tool to test functional as well as non-deterministic properties of Curry programs. CurryCheck supports unit tests and tests of I/O operations with fixed inputs as well as property tests which are parameterized over some arguments. In the latter case, they are executed with test inputs which are systematically generated for the given argument types. Moreover, CurryCheck also supports specification and contract testing if such information is present in the source program.

To simplify and, thus, enhance the use property testing, properties can be interspersed in the source program and are automatically extracted by CurryCheck. Hence, CurryCheck supports test-driven program development methods like extreme programming. Properties are not only useful to obtain more reliable programs, but they can also be used by automated documentation tools to describe the intended meaning of operations, a feature which has been recently added to the CurryDoc [21] documentation tool.[5]

For future work we plan to extend the functionality of CurryCheck (the current version does not support the generation of floating point numbers and functional values). Furthermore, we intend to integrate into CurryCheck more features that can help to improve the reliability of the source code, like abstract interpretation to approximate specific run-time properties [18,35],

[5] See `www.informatik.uni-kiel.de/~pakcs/lib/Combinatorial.html` for an example.

or program covering to show whether the test data was sufficient to reach all parts of a source program.

Acknowledgements. The author is grateful to Jan-Patrick Baye for implementing an initial version of CurryCheck.

References

1. C. Amaral, M. Florido, and V. Santos Costa. PrologCheck - property-based testing in Prolog. In *Proc. of the 12th International Symposium on Functional and Logic Porgramming (FLOPS 2014)*, pages 1–17. Springer LNCS 8475, 2014.
2. S. Antoy, R. Echahed, and M. Hanus. A needed narrowing strategy. *Journal of the ACM*, 47(4):776–822, 2000.
3. S. Antoy and M. Hanus. Declarative programming with function patterns. In *Proceedings of the International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR'05)*, pages 6–22. Springer LNCS 3901, 2005.
4. S. Antoy and M. Hanus. Overlapping rules and logic variables in functional logic programs. In *Proceedings of the 22nd International Conference on Logic Programming (ICLP 2006)*, pages 87–101. Springer LNCS 4079, 2006.
5. S. Antoy and M. Hanus. Set functions for functional logic programming. In *Proceedings of the 11th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP'09)*, pages 73–82. ACM Press, 2009.
6. S. Antoy and M. Hanus. Functional logic programming. *Communications of the ACM*, 53(4):74–85, 2010.
7. S. Antoy and M. Hanus. New functional logic design patterns. In *Proc. of the 20th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2011)*, pages 19–34. Springer LNCS 6816, 2011.
8. S. Antoy and M. Hanus. Contracts and specifications for functional logic programming. In *Proc. of the 14th International Symposium on Practical Aspects of Declarative Languages (PADL 2012)*, pages 33–47. Springer LNCS 7149, 2012.
9. S. Antoy and M. Hanus. Default rules for Curry. In *Proc. of the 18th International Symposium on Practical Aspects of Declarative Languages (PADL 2016)*, pages 65–82. Springer LNCS 9585, 2016.
10. S. Antoy, M. Hanus, and S. Libby. Proving non-deterministic computations in Agda. In *Proc. of the 24th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2016), to appear in EPTCS*, 2016.
11. F.L. Bauer, M. Broy, R. Gnatz, W. Hesse, B. Krieg-Brückner, H. Partsch, P. Pepper, and H. Wössner. Towards a wide spectrum language to support program specification and program development. *ACM SIGPLAN Notices*, 13(12):15–24, 1978.
12. J.P. Bernardy, P. Jansson, and K. Claessen. Testing polymorphic properties. In *Proc. 19th European Symposium on Programming Languages and Systems (ESOP2010)*, pages 125–144. Springer LNCS 6012, 2010.
13. B. Braßel, M. Hanus, and F. Huch. Encapsulating non-determinism in functional logic computations. *Journal of Functional and Logic Programming*, 2004(6), 2004.
14. B. Braßel, M. Hanus, B. Peemöller, and F. Reck. KiCS2: A new compiler from Curry to Haskell. In *Proc. of the 20th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2011)*, pages 1–18. Springer LNCS 6816, 2011.
15. J. Christiansen and S. Fischer. EasyCheck - test data for free. In *Proc. of the 9th International Symposium on Functional and Logic Programming (FLOPS 2008)*, pages 322–336. Springer LNCS 4989, 2008.

16. K. Claessen and J. Hughes. QuickCheck: A lightweight tool for random testing of Haskell programs. In *International Conference on Functional Programming (ICFP'00)*, pages 268–279. ACM Press, 2000.
17. K. Claessen and J. Hughes. Testing monadic code with QuickCheck. *ACM SIGPLAN Notices*, 37(12):47–59, 2002.
18. M. Fähndrich and F. Logozzo. Static contract checking with abstract interpretation. In *Proc. of the Conference on Formal Verification of Object-oriented Software (FoVeOOS 2010)*, pages 10–30. Springer LNCS 6528, 2011.
19. J.C. González-Moreno, M.T. Hortalá-González, F.J. López-Fraguas, and M. Rodríguez-Artalejo. An approach to declarative programming based on a rewriting logic. *Journal of Logic Programming*, 40:47–87, 1999.
20. M. Hanus. A unified computation model for functional and logic programming. In *Proc. of the 24th ACM Symposium on Principles of Programming Languages (Paris)*, pages 80–93, 1997.
21. M. Hanus. CurryDoc: A documentation tool for declarative programs. In *Proc. 11th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2002)*, pages 225–228. Research Report UDMI/18/2002/RR, University of Udine, 2002.
22. M. Hanus. Declarative processing of semistructured web data. In *Technical Communications of the 27th International Conference on Logic Programming*, volume 11, pages 198–208. Leibniz International Proceedings in Informatics (LIPIcs), 2011.
23. M. Hanus. Functional logic programming: From theory to Curry. In *Programming Logics - Essays in Memory of Harald Ganzinger*, pages 123–168. Springer LNCS 7797, 2013.
24. M. Hanus, S. Antoy, B. Braßel, M. Engelke, K. Höppner, J. Koj, P. Niederau, R. Sadre, and F. Steiner. PAKCS: The Portland Aachen Kiel Curry System. Available at `http://www.informatik.uni-kiel.de/~pakcs/`, 2016.
25. M. Hanus (ed.). Curry: An integrated functional logic language (vers. 0.9.0). Available at `http://www.curry-language.org`, 2016.
26. M. Hermenegildo. A documentation generator for (C)LP systems. In *Proc. of the 1st International Conference on Computation Logic (CL 2000)*, pages 1345–1361. Springer LNAI 1861, 2000.
27. M. Johansson, D. Rosén, N. Smallbone, and K. Claessen. Hipster: Integrating theory exploration in a proof assistant. In *Int. Conf. on Intelligent Computer Mathematics (CICM 2014)*, pages 108–122. Springer LNCS 8543, 2014.
28. P. Koopman, A. Alimarine, J. Tretmans, and R. Plasmeijer. Gast: Generic automated software testing. In *Proc. of the 14th International Workshop on Implementation of Functional Languages*, pages 84–100. Springer LNCS 2670, 2003.
29. I. Kuraj, V. Kuncak, and D. Jackson. Programming with enumerable sets of structures. In *Proc. of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 15)*, pages 37–56. ACM, 2015.
30. M. Papadakis and K. Sagonas. A PropEr integration of types and function specifications with property-based testing. In *Proc. of the 10th ACM SIGPLAN Workshop on Erlang*, pages 39–50, 2011.
31. S. Peyton Jones, editor. *Haskell 98 Language and Libraries—The Revised Report*. Cambridge University Press, 2003.
32. U.S. Reddy. Narrowing as the operational semantics of functional languages. In *Proc. IEEE Internat. Symposium on Logic Programming*, pages 138–151, Boston, 1985.
33. C. Runciman, M. Naylor, and F. Lindblad. SmallCheck and Lazy SmallCheck: automatic exhaustive testing for small values. In *Proc. of the 1st ACM SIGPLAN Symposium on Haskell*, pages 37–48. ACM Press, 2008.
34. J.R. Slagle. Automated theorem-proving for theories with simplifiers, commutativity, and associativity. *Journal of the ACM*, 21(4):622–642, 1974.
35. N. Stulova, J.F. Morales, and M. Hermenegildo. Reducing the overhead of assertion run-time checks via static analysis. In *Proc. 18th International Symposium on Principles and Practice of Declarative Programming (PPDP 2016)*. To appear, 2016.
36. P. Wadler. How to declare an imperative. *ACM Computing Surveys*, 29(3):240–263, 1997.